Analysis of TCP Fast Open

Erich Meissner
College of Computing, CS6250 Computer Networks
Georgia Institute of Technology
Atlanta, Georgia
erich.meissner@gatech.edu

***Abstract*—The purpose of this experiment is to further investigate a specific change to the Transmission Control Protocol (TCP) that reduces transfer latency. Persistent TCP connections enabled a browser to reuse a connection for multiple requests. HTTP requests that occur over *new* TCP connections can be ads that load from third party servers or images located on separate content servers. TCP Fast Open (TFO) permits sending data during the TCP handshake, ultimately reducing the load latency of a web resource by one round trip time (RTT). Three websites were fetched in this experiment, and an HTTP server (vnc4server) was run to illustrated the impact of TFO on total page load time (PLT) for different values of RTT.**

***Index Terms*—TCP Fast Open (TFO), Transmission Control Protocol (TCP), Round Trip Time (RTT), Page Load Time (PLT)**

TABLE I
HTTP://WWW.WASHINGTONPOST.COM

RTT (ms)	PLT: no TFO (s)	PLT: TFO (s)	Improv.
200	44949.512	36076.25	19.7405079726
100	5667.895	5346.968	5.66219028405
20	3843.454	3796.119	1.23157451605

TABLE II
HTTP://WWW.ADMISSION.GATECH.EDU

RTT (ms)	PLT: no TFO (s)	PLT: TFO (s)	Improv.
200	22894.077	2754.734	87.9674817203
100	3180.744	2390.503	24.8445332287
20	2102.965	1962.516	6.67861804642

I. Introduction

TCP Fast Open was enabled and disabled when downloading web pages from these four websites:

- http://www.washingtonpost.com
- http://www.admission.gatech.edu
- http://www.web.mit.edu

The raw experimental data generated for the for sites are listed in tables 1-3 respectively. These three sites were tested at RTTs of 20ms, 100ms, and 200ms. For context, an RTT of 100ms-200ms reflects a mobile device connected via WiFi and requesting a web-page, while an RTT of 20ms reflects a desktop computer connected via LAN/Ethernet.

The page load times for the gatech.edu site and web.mit.edu were substantial for both round trip times (RTTs) of 100 and 200 milliseconds. These two sites also experienced drastic improvements when TCP Fast Open (TFO) was enabled. When the RTT is only 20 milliseconds, we expect the gains to be small from enabling TFO. This expectation is because the network delay is only a small fraction of RTT when RTT is small. Therefore, the resource processing time would exceed network time. For washingtonpost.com, the improvement from enabled TFO when the RTT was 20 milliseconds is more comparable to the improvement when RTT is 100 and 200 milliseconds. However, the improvement when RTT is 20ms is just 6% for gatech.edu. The results are eye-opening here since the improvement for this site is drastically larger for RTTs of 100ms and 200ms.

II. Analysis

A. What effect does TFO have on the timing?

A general overview of the raw experimental data was given in the introduction in which we compared the results from the three sites to each other. Now we will investigate the output from each URL independently.

a) http://www.washingtonpost.com: For the Washington Post site, TFO has strong improvements for the 200ms RTT. The improvements are also as expected; the latency reduction is smallest for RTT at 20ms and largest for RTT at 200ms. The Washington Post is like the New York Time, Wall Street Journal, and Business Insider. These newspapers run several advertisements on their websites as that is the main method of generating revenue for the companies.

"Fig. 1", shows the results for the Washington Post site at a round trip time of 20 milliseconds. The improvement at this level is small as enabling TFO yields only a 1% improvement.

"Fig. 2", shows the results for the Washington Post site at a round trip time of 100 milliseconds. The improvement at this level is still small as enabling TFO yields only a 6% improvement.

TABLE III
HTTP://WWW.WEB.MIT.EDU

RTT (ms)	PLT: no TFO (s)	PLT: TFO (s)	Improv.
200	11772.987	923.812	92.1531213786
100	1158.645	521.382	55.0007120386
20	333.933	258.145	22.6955706684

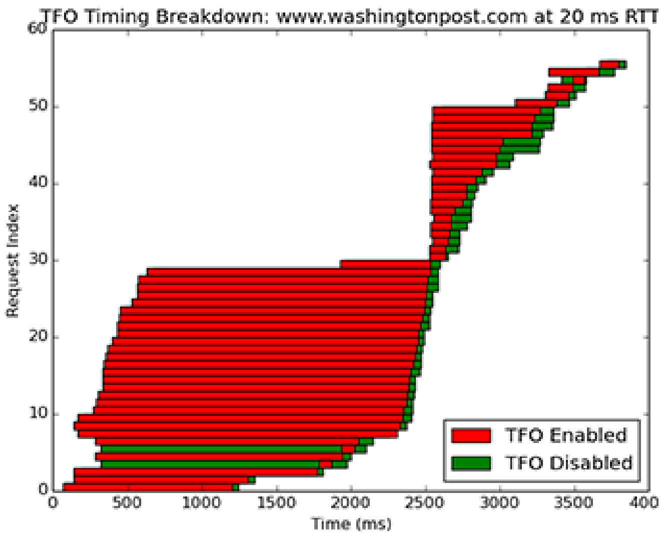

Fig. 1. washingtonpost.com site at RTT = 20ms

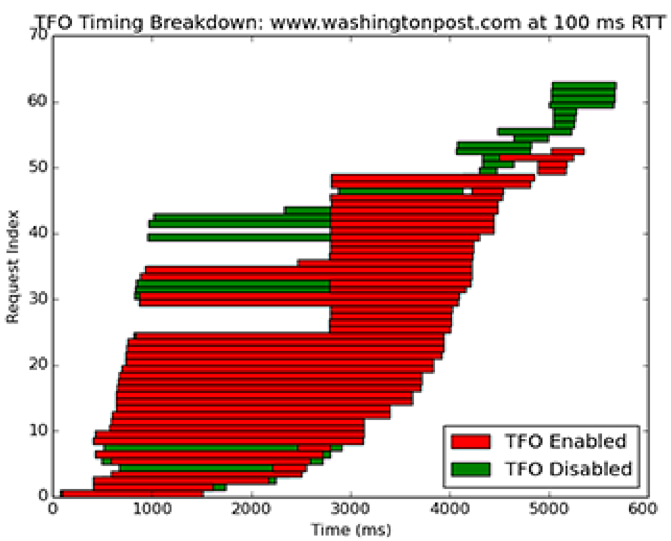

Fig. 2. washingtonpost.com site at RTT = 100ms

“Fig. 3”, shows the results for the Washington Post site a round trip time of 200 milliseconds. Now, the results are beginning to look more impressive at 20%. However, the requests at the top of the index overlap in timing for TFO enabled and disabled. The requests in question are those at index = 60. The timing, shown in milliseconds, is broad for the disabled TFO test. These results and the highest index of requests are not the same in the graphs for 20ms and 100ms RTTs for the Washington Post.

b) http://www.admission.gatech.edu: The website for Georgia Tech admissions was chosen since it is a simpler website than those like Amazon.com that can have hundreds of requests for each page. This gatech.edu site capped at about

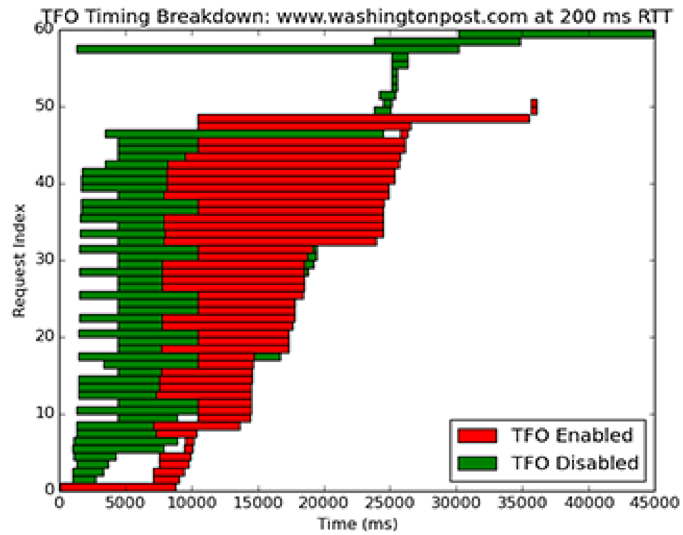

Fig. 3. washingtonpost.com site at RTT = 200ms

30 requests. Also, the site should run no ads, meaning the only other external requests may be from content servers hosting images for this web-page.

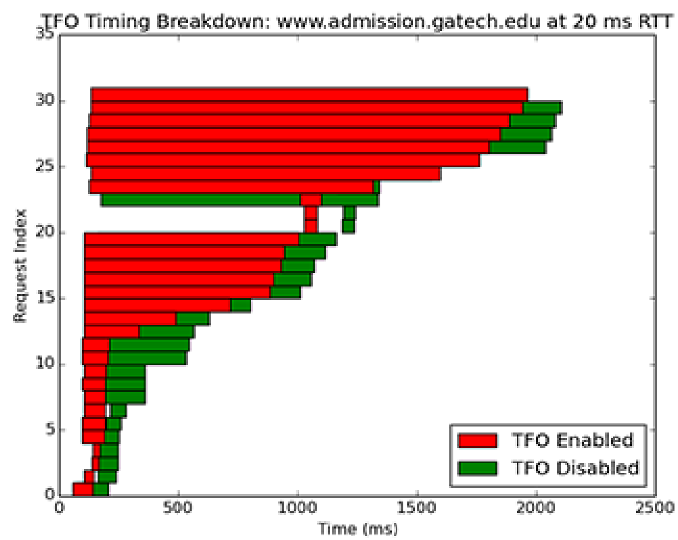

Fig. 4. gatech.edu site at RTT = 20ms

“Fig. 4”, shows the results for the Georgia Tech Admissions site at a round trip time of 20 milliseconds. The improvement at this level is small, but still better than that of the Washington Post. Enabling TFO yields a 7% improvement.

“Fig. 5”, shows the results for the Georgia Tech Admissions site at a round trip time of 100 milliseconds. The improvement at this level is strong as enabling TFO yields only a 25% improvement. This improvement is about four times greater than the 20ms RTT improvement.

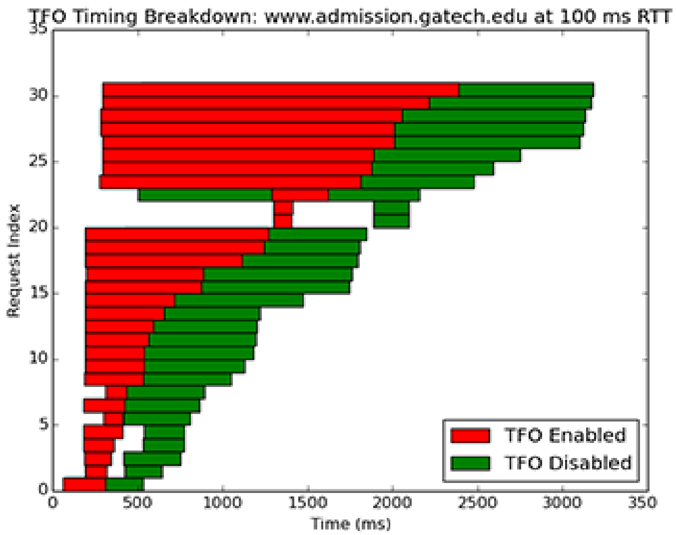

Fig. 5. gatech.edu site at RTT = 100ms

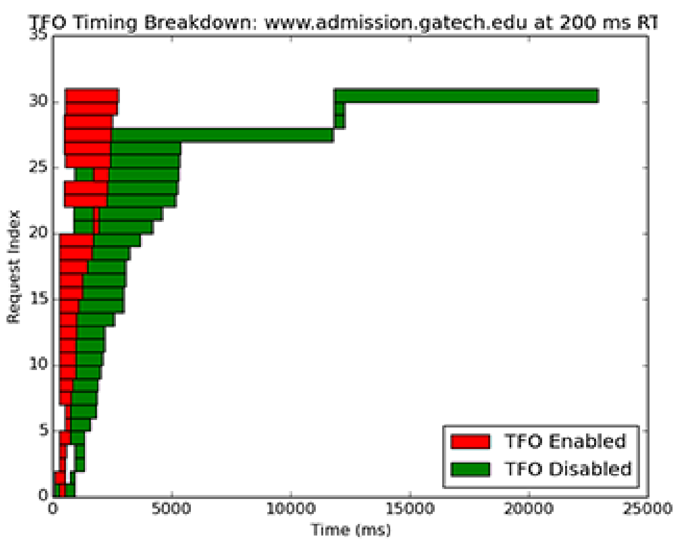

Fig. 6. gatech.edu site at RTT = 200ms

"Fig. 6", shows the results for the Georgia Tech Admissions site at a round trip time of 200 milliseconds. The improvements are beginning to be extremely impressive at 88%. An interesting observation that accounts for such a drastic improvement in applying TCP Fast Open at this 200ms RTT is the differences in timing for the highest request index. Between the 25th and 30th requested index of the admission.gatech.edu web-page, the timing is substantially smaller for TFO enabled results versus that of TFO disabled results. These requests at the top of the index surely account for the drastic improvement.

c) http://www.web.mit.edu: The homepage for the Massachusetts Institute of Technology looks similar to the admissions site of Georgia Tech; however, the number of requests for MIT's site is only six for all three RTT tests. Since the number of requests is so small, the improvement is not expected to be a large when TFO is enabled versus when it is disabled. This expectation is because TFO was built to send data during the TCP three-way handshake since so many new requests are opened for sometimes only one round trip of data sharing. If there are not many requests, then TFO cannot improve total page load time (PLT) as substantially. Given these expectations, I found the results to prove differently.

Furthermore, since the experiment is expected to perform worse on sites with HTTPS, mit.edu is a common website that does not require the secure SSL certificate. It is common due to its ranking on the Alexa top 500 sites [2].

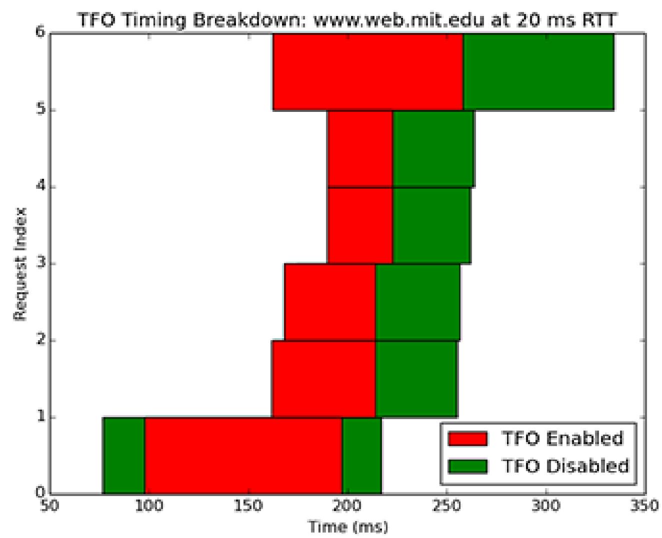

Fig. 7. web.mit.edu site at RTT = 20ms

"Fig. 7", shows the results for the Massachusetts Institute of Technology's homepage at a round trip time of 20 milliseconds. The improvement at this level is strong, and the MIT homepage proves better improvement for RTT of 20ms among all the sites in this experiment. Enabling TFO yields a 23% improvement.

"Fig. 8", shows the results for the Massachusetts Institute of Technology's homepage at a round trip time of 100 milliseconds. The improvement at this level is significant as enabling TFO yields a 55% improvement. This improvement is about twice as great as the 20ms RTT improvement.

"Fig. 9", shows the results for the Massachusetts Institute of Technology's homepage at a round trip time of 200 milliseconds. The improvements are extremely impressive at 92%. An interesting observation that accounts for such a drastic improvement in applying TCP Fast Open at this 200ms RTT is the differences in timing for the highest request index. Between the 5th and 6th requested index of the web.mit.edu web-page, the timing is substantially smaller for TFO enabled results versus that of TFO disabled results. These requests at

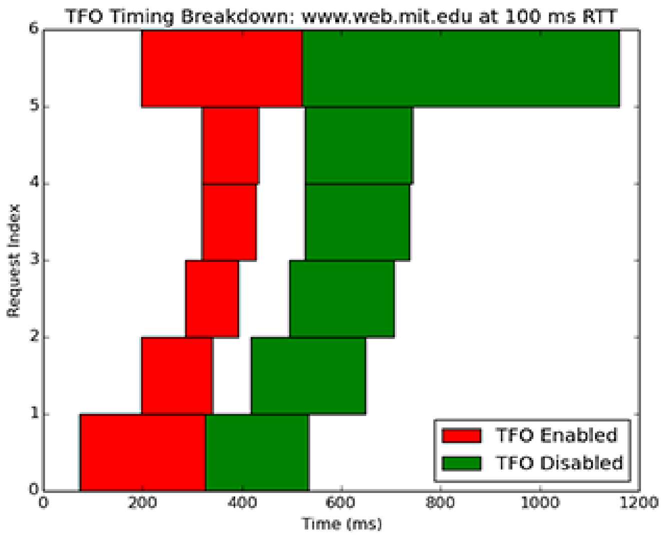

Fig. 8. web.mit.edu site at RTT = 100ms

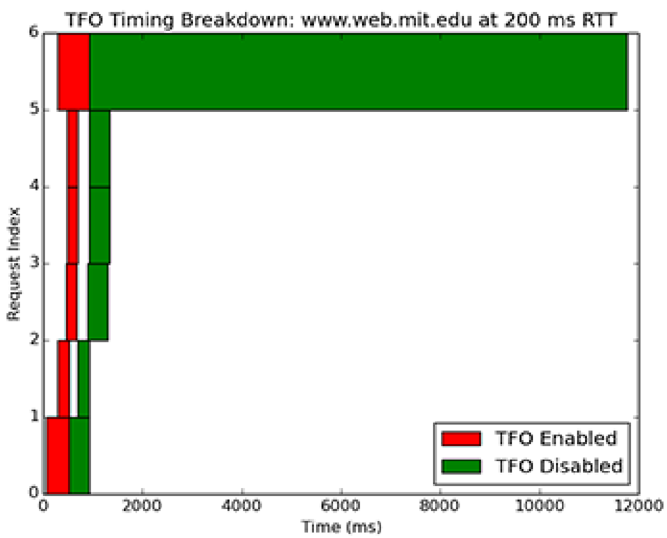

Fig. 9. web.mit.edu site at RTT = 200ms

the top of the index surely account for the drastic improvement, much like the improvement seen for the 200ms RTT gatech.edu test.

B. How does the RTT value affect these results?

TCP Fast Open (TFO) is meant to improve the total page load time (PLT) when the RTT is high. When the RTT is low, the network delay is only a small percentage of the PLT. Therefore, the resource processing time would exceed network time, meaning the improvements from TFO are smaller. This expectation proved true for all three websites.

a) http://www.washingtonpost.com: The Washington Post website yielded better improvements as the RTT values increased. These results were expected.

b) http://www.admission.gatech.edu: The admissiong site of Georgia Tech yielded better improvements as the RTT values increased. These results were also expected.

c) http://www.web.mit.edu: Finally, the homepage for the Massachusetts Institute of Technology yielded better improvements as the RTT values increased. The expectation that TFO improvements rise with higher RTTs proved true for all three sites.

C. Does the particular content available at this URL lend itself to performance enhancements provided by TFO?

Websites that are simple, which can be those with lower than 10 requests and not much dense content like embedded images, should not see substantial improvements from TCP Fast Open. The Washington Post has many advertisements from third parties running on their site as well as plenty of images. The expectation was that this site would yield weaker improvements than that of simpler sites. The resource processing time on a site like the Washington Post exceeds the network time, meaning smaller improvements from TFO being enabled. Conversely, the two simple sites, web.mit.edu and admission.gatech.edu, can also expect just as significant improvements from TFO. For simple sites like these two, the Chrome browser spends most of its time waiting for network transfers rather than processing the retrieved content. The greatest improvements should be seen at gatech.edu and mit.edu, therefore.

a) http://www.washingtonpost.com: The Washington Post website yielded the weakest improvement at all three levels, yet TFO still did accelerate total PLT by 20% at an RTT of 200ms.

b) http://www.admission.gatech.edu: The admission site of Georgia Tech yielded drastic improvements as expected for a site of this content.

c) http://www.web.mit.edu: Finally, the homepage for the Massachusetts Institute of Technology also yielded strong results for a site with simple content and no advertisements.

D. Were the results surprising in any way?

The results proved to be inline with expectations based on understanding of TCP Fast Open and how it accelerates page load times. Overall, the only surprise was the extent to which it improved the load times.

a) http://www.washingtonpost.com: The Washington Post website yielded the weakest improvement at all three levels; however, the improvements were almost identical to that of the paper from researchers at Google and Berkeley [1]. Their paper investigated improvements from the Wall Street Journal and New York Times, which are both sites quite similar to The Washington Post. This experiment yielded improvements almost identical to that of the two news sites studied in the TFO paper.

b) http://www.admission.gatech.edu: The admission site of Georgia Tech yielded drastic improvements as expected for a site of this content. The surprise is that this site improved by 88% for its total PLT at an RTT of 200ms. The TFO paper generated 41% improvement from enabling TFO for Wikipedia, and similar results were expected for gatech.edu [1]. However, the improvements were much larger.

c) http://www.web.mit.edu: Finally, the homepage for the Massachusetts Institute of Technology also yielded strong results for a site with simple content and no advertisements. The surprise is again that this site improved by 92% for its total PLT at an RTT of 200ms. The TFO paper generated 41% improvement from enabling TFO for Wikipedia, and similar results were expected for gatech.edu [1]. For both the mit.edu site and gatech.edu site, the results were more than double what the TFO paper found for Wikipedia.

E. Include relevant graphs from the output-figures folder.

Nine graphs are included in this analysis. The round trip times (RTTs) of 20ms, 100ms, and 200ms for each of the three websites are found inline with comments on the figures.

III. Brief Summary of Findings and Conclusions Drawn from Experimental Results

Overall, the experiment supplied results that agreed well with expectation set from the original TCP Fast Open study on sites like nytimes.com, wsj.com, wikipedia.com, and amazon.com [1]. A reason for why the two simple sites in this study, admission.gatech.edu and web.mit.edu, performed so well is that they exemplify the perfect sites that can benefit from a concept like TFO. For these webpages, the browser is not waiting on the retrieved resources to be processed. Rather, these two sites are spending most time waiting for network transfers. For this reason, these sites both improve about 90% in total page load time for an RTT of 200ms.

IV. Which Website(s) have the Best and Worst Performance of TCP Fast Open over TCP.

For similar reasons outlined above, TCP Fast Open improves page load time the most for sites that wait longer on network transfer times rather than waiting on resources to process. Websites that are not heavy on content will see the best improvements from applying TCP Fast Open over TCP. Those that are heavy on content and have resource processing time exceed network time with find the worst performance of TCP Fast Open over TCP; however, they will find an improvement nonetheless.

References

[1] Sivasankar Radhakrishnan, Yuchung Cheng, Jerry Chu, Arvind Jain, “TCP Fast Open” Google, ICSI University of California, Berkeley, April 2011.

[2] Alexa Top 500 Global Sites, http://www.alexa.com/topsites.